#17

PATENT
Attorney Docket No. 99-829CPA1

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

7/7/04

In re Application of:	)	
	)	
Himanshu S. SINHA	)	Group Art Unit: 2142
	)	
Serial No.: 09/425,088	)	Examiner: D. Blair
	)	
Filed: October 22, 1999	)	
	)	
For: SERVICE LEVEL AGREEMENTS AND	)	
MANAGEMENT THEREOF	)	

RECEIVED

JUN 3 0 2004

Technology Center 2100

APPEAL BRIEF

Mail Stop Appeal Brief – Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

This Appeal Brief is submitted in response to the Final Office Action mailed January

28, 2004 and in support of the Notice of Appeal filed April 23, 2004.

I.   REAL PARTY IN INTEREST

The real party in interest in this appeal is Verizon Corporate Services Group Inc.

II.   RELATED APPEALS AND INTERFERENCES

Appellant is unaware of any related appeals and interferences.

III.   STATUS OF CLAIMS

Claims 1-5 and 7-17 are pending in this application. All of the pending claims are the

subject of the present appeal.

IV.   STATUS OF AMENDMENTS

No Amendment has been filed subsequent to the Final Office Action mailed January 28,

2004.

V.   SUMMARY OF THE INVENTION

The present invention is directed to service level agreements and managing service level

agreements (SLAs) between composite service providers, such as client 100, and component

service providers, such as service provider 120 (specification – page 1, lines 4-18 and page 8,

lines 1-12 and Fig. 2). As further described in Appellant's specification at page 8, lines 15-

17 and with reference to Fig. 2, SLA manager 110 may include an SLA admission controller

113, a performance measurement module 111 and an SLA specification module 112.   In an

exemplary embodiment, SLA manager 110 makes an initial estimate of system capacity by

measuring system performance under a simulated load. SLA manager 110 then uses

performance measurement module 111 to measure actual performance of one or more service

implementations to refine the initial capacity estimate (specification – page 11, lines 13-19). In

this manner, SLA manager 110 can improve capacity estimates generated via simulations using

actual performance results, which more accurately reflect client access patterns (specification –

page 12, lines 4-12).

As further described at page 12, line 13 to page 13, line 7, a capacity of a service provider 120 may be denoted by a number of tokens. Each client organization, such as client 100a, may be assigned tokens associated with interacting with a service provider 120 via SLA manager 110. For example, admission controller 113 may receive a request R_i from client 100a and check whether a sufficient number of tokens exist in client 100a's account. If so, request R_i is accepted and the number of tokens needed for R_i is deducted from client 100a's account. When request R_i is completed, the number of tokens that were deducted are credited back to client 100a's account.

VI.   ISSUES

Whether claims 1-5, 7-10 and 12-14 are patentable under 35 U.S.C. § 103(a) over Somers (U.S. Patent No. 6,243,396) in view of Tunnicliffe et al. (U.S. Patent No. 6,272,110; hereinafter Tunnicliffe);

Whether claims 11 and 15 are patentable under 35 U.S.C. § 103(a) over Somers in view of Tunnicliffe and further in view of Ball et al. (U.S. Patent No. 6,446,200; hereinafter Ball); and

Whether claims 16 and 17 are patentable under 35 U.S.C. § 103 over Somers in view of Tunnicliffe and further in view of Aronberg et al. (U.S. Patent No. 6,117,188; hereinafter Aronberg) and further in view of Knight et al. (U.S. Patent No. 6,442,608; hereinafter Knight).

VII.   GROUPING OF CLAIMS

Appellant is satisfied to let claims 1-5, 7-10 and 12-14 stand or fall together. Appellant

is also satisfied to let claims 11 and 15 stand or fall together. Each of claims 16 and 17 do not

stand or fall together with any of the other claims for the reasons discussed in the Argument

section below.

VIII.   ARGUMENT

A.   The Rejections

1. Claims 1-5, 7-10 and 12-14 are patentable over Somers in view of Tunnicliffe.

The initial burden of establishing a *prima facie* basis to deny patentability to a claimed

invention always rests upon the Examiner. In re Oetiker, 977 F.2d 1443, 24 USPQ2d 1443

(Fed. Cir. 1992). In rejecting a claim under 35 U.S.C. § 103, the Examiner must provide a

factual basis to support the conclusion of obviousness. In re Warner, 379 F.2d 1011, 154

USPQ 173 (CCPA 1967). Based upon the objective evidence of record, the Examiner is

required to make the factual inquiries mandated by Graham v. John Deere Co., 86 S.Ct. 684,

383 U.S. 1, 148 USPQ 459 (1966). The Examiner is also required to explain how and why

one having ordinary skill in the art would have been realistically motivated to modify an

applied reference and/or combine applied references to arrive at the claimed invention.

Uniroyal, Inc. v. Rudkin-Wiley Corp., 837 F.2d 1044, 5 USPQ2d 1434 (Fed. Cir. 1988).

In establishing the requisite motivation, it has been consistently held that the requisite

motivation to support the conclusion of obviousness is not an abstract concept, but must stem

from the prior art as a whole to impel one having ordinary skill in the art to modify a reference

or to combine references with a reasonable expectation of successfully achieving some

particular realistic objective. See, for example, Interconnect Planning Corp. v. Feil, 227

USPQ 543 (Fed. Cir. 1985). Consistent legal precedent admonishes against the indiscriminate

combination of prior art references. Carella v. Starlight Archery, 804 F.2d 135, 231 USPQ

644 (Fed. Cir. 1986); Ashland Oil, Inc. v. Delta Resins & Refractories, Inc., 776 F.2d 281,

227 USPQ 657 (Fed. Cir. 1985).

With these principles in mind, the arguments below use claim 1 as representative of the

group of claims including claims 1-5, 7-10 and 12-14. Claim 1 recites an apparatus that

includes a service level agreement manager disposed between a client computer system and a

service implementation. The service level agreement manager includes an admission

controller, a performance measurement module and a specification module. Claim 1 recites

that the performance measurement module is configured to measure performance of the service

implementation and modify an estimated capacity of the service provider based on the

measured performance.

The Final Office Action points to the rejection made in the previous Office Action

dated September 12, 2003 for the detailed rejection (Final Office Action – page 2). The Office

Action dated September 12, 2003 (hereinafter "Office Action") states that Somers discloses a

performance measurement module that is configured to measure performance of a service

implementation, but admits that Somers does not disclose modifying an estimated capacity of

the service provider based on the measured performance (Office Action – page 3). The Office

Action, however, states that Tunnicliffe discloses this feature and points to col. 6, lines 53-67

and col. 7, lines 1-3 of Tunnicliffe for support (Office Action – page 3). Appellant

respectfully disagrees.

Tunnicliffe is directed to a method and apparatus for managing at least part of a

communications network (Tunnicliffe – col. 1, lines 7-9). Tunnicliffe at col. 6, line 37 to col.

7, line 3 discloses that agents may negotiate service level agreements (SLAs) with each other

through SLA proposals. The negotiation model is implemented using a declarative knowledge

base (KB) and a procedural KB. Tunnicliffe also discloses that the agents can negotiate over

the price of a service by generating a price offer and continuing to counter-propose the initial

offer (Tunnicliffe – col. 6, line 59 to col. 7, line 2). This portion of Tunnicliffe clearly does

not disclose measuring a performance of the service implementation and modifying an

estimated capacity of the service provider based on the measured performance, as recited in

claim 1. In contrast, this portion of Tunnicliffe merely discloses that agents may use

knowledge bases to negotiate SLAs, such as when negotiating an SLA based on price.

Tunnicliffe does disclose that a network operator monitors network traffic entering the

network from each of customer sites 203 at points 207, 208 and 209 and predicts further

demand (col. 4, lines 10-23). The network operator may pro-actively reconfigure resources to

cope with increased traffic, detect when customers are likely to exceed agreed levels of

service, agree to carry the excess traffic, cut back on inter-site traffic or block excess traffic

(Tunnicliffe – col. 4, lines 24-40). Tunnicliffe also discloses that predictor 13 (Fig. 1) may be

used to predict when the demand will exceed the capacity, the extent of the predicted excess

and when the demand will return below capacity (Tunnicliffe – col. 4, lines 40-59 and Fig. 3).

In summary, these portions of Tunnicliffe may disclose that a network operator may predict short term demand on the network and make various adjustments based on the amount of demand. None of these portions of Tunnicliffe, however, discloses modifying an estimated capacity of the service provider based on the measured performance, as recited in claim 1. In other words, determining when short term demand may exceed agreed levels of service and making any of the adjustments discussed in Tunnicliffe is not equivalent to and does not suggest modifying an estimated capacity of a service provider based on a measured performance, as recited in claim 1.

Therefore, the portions of Tunnicliffe pointed to on page 3 of the Office Action, as well as the other portions discussed above relating to predicting short term demand on the network, do not disclose or suggest modifying an estimated capacity of a service provider based on the measured performance, as recited in claim 1.

In the Response to Arguments section, the Final Office Action states that "the prediction of short-term demand is considered a modification of estimated capacity based on measured performance" (Final Office Action – page 3). The Final Office Action adds that there "is nothing in the claim language that states that the performance measurement module is not in communication with an operator nor is there any claim language that states that the estimated capacity is not short-term" (Final Office Action – pages 3-4).

It appears that the Examiner is equating demand on a network with capacity. Appellant strongly disagrees with this analysis. Demand on a network is a measure of a load on the

network, such as the sum of all traffic being transmitted on the network over some period of

time. Capacity, in contrast, is a measure of how much traffic can be handled. Demand and

capacity, based on both their normal meaning and the way that they are typically used in this

art, are clearly not equivalent. Therefore, the mere fact that Tunnicliffe discloses determining

a short-term demand on a network does not mean that Tunnicliffe discloses or suggest

modifying an estimated capacity of the service provider based on the measured performance, as

recited in claim 1.

Therefore, as a factual matter, the combination of Somers and Tunnicliffe does not

disclose or suggest each of the features of claim 1.

In addition, even if, for the sake of argument, the combination of Somers and

Tunnicliffe could reasonably be construed to disclose each of the features of claim 1, the Final

Office Action does not provide the requisite motivation under 35 U.S.C. § 103 as to why it

would have been obvious to combine these references.

For example, the Office Action states that it would have been obvious to combine "the

teachings of Somers regarding a service level agreement implementation with the teachings of

Tunnicliffe regarding modifying an estimated system capacity based on the measured

performance because changing an estimated capacity provides more flexibility for clients

(Tunnicliffe col. 1, lines 11-35)" (Office Action – page 3). Appellant respectfully disagrees.

Initially, Appellant notes that the alleged motivation (i.e., to provide more flexibility for

clients) is merely a conclusory statement regarding an alleged benefit of the combination. Such

motivation does not satisfy the requirements of 35 U.S.C. § 103. In this respect, Appellant

relies upon In re Deuel, 51 F.3d 1552, 34 USPQ2d 1210 (Fed. Cir. 1995), wherein it was held

that generalizations do not establish the realistic motivation to modify a specific reference in a

specific manner to arrive at a specifically claimed invention.

Secondly, the portion of Tunnicliffe pointed to in the Office Action (i.e., col. 1, lines

11-35) also provides no support for the combination. Tunnicliffe at col. 1, lines 11-35

references the Description of the Prior Art section of Tunnicliffe. This section discloses that

SLAs typically specify bandwidth levels that the customer is allowed to use and other quality

factors. Tunnicliffe discloses that is difficult for customers to predict bandwidth requirements

and penalty payments may be involved by exceeding the allocated level (Tunnicliffe – col. 1,

lines 11-24). This section of Tunnicliffe also discloses that if agreed bandwidth level is

exceeded by the customer, the service provider has to decide whether to allocate extra

resources to the customer, determine whether this is possible and how it can be achieved. This

involves analysis of the network to see what bandwidth is available (Tunnicliffe – col. 1, lines

25-35). This portion of Tunnicliffe provides no suggestion for combining Tunnicliffe, which

is directed to a method and apparatus for managing at least part of a network, with Somers,

which is directed to a management system having a hierarchical structure of management

authorities controlling network resources. The mere fact that one reference allegedly provides

some missing disclosure with respect to a claim does not satisfy the requirements of 35 U.S.C.

§ 103 as to why it would have been obvious to combine the references. Appellant asserts that

it would not have been obvious to combine these two references without the benefit of

Appellant's disclosure. Moreover, the failure of the Office Action to logically explain why

one skilled in the art would have realistically been motivated to combine the teachings of

Somers and Tunnicliffe further supports Appellant's assertion.

Therefore, Appellant respectfully submits that the imposed rejection of claims 1-5, 7-

10 and 12-14 under 35 U.S.C. § 103 for obviousness based on the combination of Somers and

Tunnicliffe is improper. Accordingly, reversal of the rejection is respectfully requested.

2. Claims 11 and 15 are patentable under 35 U.S.C. § 103 over Somers in view of

Tunnicliffe and in further view of Ball.

The arguments below use claim 11 as representative of the group of claims including

claims 11 and 15. Claim 11 is dependent on claim 8 and recites that the invocation

infrastructure comprises Hypertext Transport Protocol. The Final Office Action once again

points to the rejection made in the Office Action dated September 12, 2003 for the detailed

rejection.

The Office Action admits that the combination of Somers and Tunnicliffe does not

disclose this feature, but states that Ball discloses a network wherein the invocation

infrastructure comprises http and points to col. 8, lines 1-24 for support (Office Action – page

10). The Office Action also states that it would have been obvious to combine the teachings of

Somers-Tunnicliffe with the teachings of Ball regarding the use of http because "the use of

http reflects the clients interaction with a service system" (Office Action – page 10). Appellant

respectfully disagrees.

Ball is directed to a system for collecting and aggregating data from network entities for

data consuming applications (Ball – Abstract). Ball at col. 8, lines 1-24 discloses that a data flow 130 may be initiated by a user 131 making a call to a remote access concentrator (RAC) 132. RAC 132 may send a remote authentication dial-in user service (RADIUS) network accounting record (NAR) to accounting process 14 and the remote user may check email, look at a web server and transfer a file. For each transaction, the accounting process 14 captures the IP traffic, generates email, http, and ftp network accounting records. This portion of Ball merely discloses that an accounting process 14 may capture hypertext transport protocol accounting records. This is not equivalent to an invocation infrastructure that communicates between a plurality of client processes and a plurality of service level managers, where the invocation infrastructure comprises hypertext transport protocol, as recited in claim 11.

Appellant also notes that the alleged motivation (i.e., because the use of http reflects the clients interactions with a service system) is merely a conclusory statement regarding http. Such motivation does not satisfy the requirements of 35 U.S.C. § 103. In re Deuel, supra. It is apparent that the Examiner's approach to the ultimate legal conclusion of obviousness under 35 U.S.C.§ 103 amounts to a retrospective assessment as to how the claimed invention works and then combining unrelated references to arrive at the claimed invention. This type of reverse engineering approach to the obviousness issue under 35 U.S.C. § 103 has been repeatedly judicially condemned. Uniroyal, Inc. v. Rudkin-Wiley Corp., supra; Panduit Corp. v. Dennison Mfg. Co., supra. Absent such hindsight reasoning, one of ordinary skill in the art would not have been motivated to combine the references in the manner suggested by the Examiner.

Therefore, Appellant respectfully submits that the imposed rejection of claims 11 and 15 under 35 U.S.C. § 103 for obviousness based on the combination of Somers, Tunnicliffe and Ball is improper. Accordingly, reversal of the rejection is respectfully requested.

3. Claims 16 and 17 are patentable under 35 U.S.C. § 103 over Somers in view of Tunnicliffe in further view of Aronberg and in further view of Knight.

a) Claim 16

Claim 16 is dependent on claim 8 and recites that each of the plurality of client processes is assigned a number of tokens and when determining whether to accept the request from a first client process to a first service level manager, the first service level manager is configured to determine whether to accept the request based on the number of tokens associated with the first client process.

As to claim 16, the Final Office Action refers to the Office Action dated September 12, 2003 for the detailed rejection. The Office Action admits that Somers and Tunnicliffe do not disclose the features of claim 16 and states that Knight discloses a network in which each of a number of client processes is assigned a number of sessions and that a first service level manager is configured to determine whether to accept a request from a first client process based on the number of sessions associated with the first client process and points to col. 23, line 33 to col. 25, line 48 for support (Office Action – pages 10-11). The Office Action admits that Knight does not disclose the use of tokens associated with a client process, but states that Aronberg discloses the use of a fixed number of tokens used to regulate network access and points to col. 4, line 56 to col. 5, line 30 for support (Office Action – page 11). The Office Action further states that it

- 12 -

would have been obvious to combine Knight and Aronberg "because tokens provide a functional

alternative to the counter as implemented in Knight" (Office Action – page 11). Appellant

respectfully disagrees.

As admitted in the Office Action, Knight does not disclose assigning a number of tokens

to each of a number of client processes, as recited in claim 16, but indicates that tokens may be

used as a functional alternative to the counter used in Knight. The Office Action also relies upon

Aronberg as disclosing the use of tokens (Office Action – page 11). More particularly, Aronberg

at col. 4, line 56 to col. 5, line 30 has been alleged to disclose the use of tokens to regulate

network access. This portion of Aronberg actually discloses that a dialog box 401 may be used

by an administrator to control the number of concurrent software distributions (See Fig. 4). A

token server box 401B may be checked to indicate use of the token server feature. An agent may

then wait for a token to allow the user to download applications (Aronberg – col. 5, lines 1-15).

This portion of Aronberg does not <u>disclose determining whether to accept a request based on the</u>

<u>number of tokens associated with a client process</u>, as recited in claim 16.

In response to some of these arguments made in a previous response, the Response to

Arguments section of the Office Action states that "col. 5, lines 7-13 of Aronberg describe

allotting a number of tokens to users that access the network some form of client process (in

Aronberg, col. 2, lines 40-61, the users access a server thus using some sort of client process.)"

(Office Action – page 12).

Aronberg at col. 5, lines 7-13 discloses that field 401E of Fig. 4 allows the administrator

to enter the maximum number of simultaneous users permitted to download an application. This

value represents the maximum number of available tokens. Aronberg at col. 2, lines 40-61

discloses that a server controls the number of users simultaneously distributing software by using

a token server. The token server issues a token authorization signal to allow an agent

workstation to download the software. Neither of these portions of Aronberg discloses or

suggests determining whether to accept a request based on the number of tokens associated with

a client process, as recited in claim 16.

Therefore, even if Aronberg and Knight were combined with the combination of Somers

and Tunnicliffe, the claimed invention would not result.

In addition, even if, for the sake of argument, the disclosures of Somers, Tunnicliffe,

Knight and Aronberg could reasonably be construed to disclose each of the features of claim

16, the Office Action does not provide the requisite motivation under 35 U.S.C. § 103 as to

why it would have been obvious to combine these references.

For example, the Office Action states that it would have been obvious to combine the

teachings of Somers-Tunnicliffe with the teachings of Knight-Aronberg "because limiting access

of specific clients would ensure a more consistent level of service for all clients" (Office Action

– page 11). Appellant respectfully disagrees.

Initially, Appellant notes that no portion of any of the four references is pointed to as

providing objective motivation for combining the disclosures of Aronberg and Knight with the

combination of Somers and Tunnicliffe. The motivation provided in the Office Action (i.e., to

ensure a more consistent level of service for all clients) is merely a conclusory statement

regarding an alleged benefit resulting from the combination. Such motivation does not satisfy

the requirements of 35 U.S.C. § 103. In re Deuel, supra.

In addition, Knight is directed to managing the access of a network system using

distributed authorization controlled by distributed nodes (Knight – col. 1, lines 15-18).

Aronberg, in contrast, is directed a system for distributing software in a network environment

(Aronberg – Abstract). Knight and Aronberg are clearly directed to different environments and

Appellant asserts that it would not have been obvious to combine features from these disparate

environments without the benefit of Appellant's disclosure.

Once again, it is apparent that the Examiner's approach to the ultimate legal

conclusion of obviousness under 35 U.S.C.§ 103 amounts to a retrospective assessment as to

how the claimed invention works and then combining unrelated references to arrive at the

claimed invention. This type of reverse engineering approach to the obviousness issue under

35 U.S.C. § 103 has been repeatedly judicially condemned. Uniroyal, Inc. v. Rudkin-Wiley

Corp., supra; Panduit Corp. v. Dennison Mfg. Co., supra. Absent such hindsight reasoning,

one of ordinary skill in the art would not have been motivated to combine the references in the

manner suggested by the Examiner.

Therefore, Appellant respectfully submits that the imposed rejection of claim 16 under

35 U.S.C. § 103 for obviousness based on the combination of Somers, Tunnicliffe, Aronberg

and Knight is improper. Accordingly, reversal of the rejection is respectfully requested.

b) Claim 17

Claim 17 is dependent on claim 16 and recites that when the request from the first client

process is accepted, the first service level manager is further configured to deduct a number of

tokens from the first client process. The Final Office Action points to the Office Action dated

September 12, 2003 for the detailed rejection (Final Office Action – page 3). The Office Action states that Knight discloses deducting a count associated with a first client process when a request from a client process is accepted and points to Knight at col. 23, line 33 to col. 25, line 48 for support (Office Action – page 11). The Office Action further states that it would have been obvious to use tokens instead of a count for the reasons discussed with respect to claim 16. Appellant respectfully disagrees.

Knight discloses that an entity, such as a company, may be assigned a threshold value associated with a maximum number of sessions that may be established for the entity at a particular time (Knight – col. 22, lines 35-47). Assigning sessions to an entity, such as a company, is not equivalent to assigning sessions to each of a plurality of client processes. In addition, Knight does not further disclose deducting a number of sessions from the client process if the request is accepted, as required by claim 17. Rather, Knight, as best understood by Appellant, merely compares the local session threshold value with the local session counter value to determine whether to authorize the request (Knight – col. 23, line 45 to col. 24, line 65). This is not equivalent to and does not suggest deducting a number of sessions or tokens from a first client process if the request is accepted. Aronberg also does not disclose deducting a number of tokens from the first client process when a request from the first client process is accepted. Therefore, the combination of Somers, Tunnicliffe, Aronberg and Knight does not disclose or suggest this feature of claim 17.

In response to some of these arguments made in a previous response, the Response to Arguments section of the Office Action states that Knight discloses that deducting a number of sessions from a particular entity also deducts the number of sessions for the associated client

process used to access the network and points to Fig. 1A of Knight for support (Office Action –

page 13). Appellant respectfully disagrees.

Fig. 1A of Knight illustrates a system 100 including one or more clients 102a-d, one or

more network access servers 104 and 106, one or more local distributed session counters (DSCs)

108 and 110, a central or authoritative DSC 112 and a network 114 (Knight – col. 7, lines 13-18).

This portion of Knight, or any other portion, does not disclose or suggest that any of DSCs 108-

112 deduct tokens (or sessions) from a first client process when the first client process is

accepted, as recited in claim 17.

Therefore, even if Aronberg and Knight were combined with the combination of Somers

and Tunnicliffe, the claimed invention would not result.

In addition, Appellant respectfully submits that the motivation to combine Somers,

Tunnicliffe, Aronberg and Knight does not satisfy the requirements of 35 U.S.C. § 103 for the

reasons given above with respect to claim 16.

Therefore, Appellant respectfully submits that the imposed rejection of claim 17 under

35 U.S.C. § 103 for obviousness based on the combination of Somers, Tunnicliffe, Aronberg

and Knight is improper. Accordingly, reversal of the rejection is respectfully requested.

IX.   CONCLUSION

In view of the foregoing arguments, Appellant respectfully solicits the Honorable Board

to reverse the Examiner's rejections of claims 1-5 and 7-17 under 35 U.S.C. § 103.

To the extent necessary, a petition for an extension of time under 37 C.F.R. § 1.136 is

hereby made. Please charge any shortage in fees due in connection with the filing of this

paper, including extension of time fees, to Deposit Account 07-2347 and please credit any

excess fees to such deposit account.

Respectfully submitted,

By: _____
     Joel Wall
     Reg. No. 25,648

Date: June 22, 2004

Verizon Corporate Services Group Inc.
600 Hidden Ridge Drive
Mail Code HQE03H14
Irving, Texas 75038
(972) 718-4800

CUSTOMER NO. 32127

APPENDIX

1. In a system having a client computer system and a service provider computer system programmed with a service implementation, an apparatus comprising:

a service level agreement manager disposed between the client computer system and the service implementation, the service level agreement manager comprising:

an admission controller configured to control admission of the client computer system to the service implementation using a service level agreement;

a performance measurement module in communication with the admission controller and configured to:

measure performance of the service implementation, and

modify an estimated capacity of the service provider based on the measured performance; and

a specification module in communication with the admission controller and with the performance measurement module.

2. The apparatus of Claim 1, wherein the specification module is configured to compare service implementation performance data and client usage information.

3. A method for service level formation, comprising:

providing a service level agreement manager, the service level agreement manager having an admission controller, a specification module and a performance measurement

module;

establishing communication between a client computer system and the service level

agreement manager;

invoking the specification module of the service level agreement manager;

obtaining performance information from the performance measurement module;

obtaining usage information from the client;

comparing the performance information and the usage information to determine if there

exists a basis for forming a service level agreement;

measuring, by the performance measurement module, actual performance associated

with at least one service level agreement; and

modifying estimated capacity, by the service level agreement manager, based on the

measured actual performance.

4. The method of Claim 3, further comprising:

forming the service level agreement; and

providing the admission controller with specification information from the service level

agreement formed.

5. A method for managing system performance, comprising:

providing a service level agreement manager;

forming a service level agreement between a client organization and a service

organization;

receiving a request from the client organization to the service level agreement manager;

with the service level agreement manager,

determining if the request is within the scope of the service level agreement;

if the request is within the scope of the service level agreement, providing the request to a performance measurement module and to the service organization;

obtaining a result from the service organization in response to the request;

taking at least one performance measurement associated with performance response of the service organization to the request; and

checking the at least one performance measurement taken against the service level agreement;

recording the at least one performance measurement; and

modifying an estimated capacity associated with the service organization based on the at least one performance measurement.

7. The method of Claim 5, further comprising providing the result obtained to the client.

8. A network, comprising:

a plurality of service level managers;

at least one invocation infrastructure for communication between a plurality of client

processes and the plurality of service level managers; and

each service level manager of the service level managers in communication with a

respective service implementation and configured to:

receive a request from at least one of the client processes,

determine whether to accept the request based on an estimated capacity of a

service provider,

accept the request when the estimated capacity is adequate,

measure performance associated with fulfilling the request, and

modify the estimated capacity based on the measured performance.

9. The network of Claim 8, wherein the invocation infrastructure comprises a

Common Object Request Broker Architecture.

10. The network of Claim 8, wherein the invocation infrastructure comprises Java

Remote Method Invocation.

11. The network of Claim 8, wherein the invocation infrastructure comprises

Hypertext Transport Protocol.

12. A network, comprising:

a first plurality of service level managers;

at least one invocation infrastructure for communication between said first plurality of service level managers and a client process;

each service level manager of said first plurality of service level managers in communication with a respective service implementation of a first plurality of service implementations;

each said service implementation of said first plurality of service implementations in communication with at least one service level manager of a second plurality of service level managers; and

each service level manager of said second plurality of service level managers in communication with a respective service implementation of a second plurality of service level implementations, wherein at least one of the first plurality and second plurality of service level managers is configured to:

enter into a service level agreement with the client process,

receive a request from the client process,

determine whether to accept the request based on an estimated capacity of a service provider,

accept the request when the estimated capacity is adequate,

measure performance associated with fulfilling the request, and

modify the estimated capacity based on the measured performance.

13. The network of Claim 12, wherein the invocation infrastructure comprises a Common Object Request Broker Architecture.

14. The network of Claim 12, wherein the invocation infrastructure comprises Java Remote Method Invocation.

15. The network of Claim 12, wherein the invocation infrastructure comprises Hypertext Transport Protocol.

16. The network of claim 8, wherein each of the plurality of client processes is assigned a number of tokens and when determining whether to accept the request from a first client process to a first service level manager, the first service level manager is further configured to:

determine whether to accept the request based on the number of tokens associated with the first client process.

17. The network of claim 16, wherein when the request from the first client process is accepted, the first service level manager is further configured to:

deduct a number of tokens from the first client process.